

(19)



(11)

**EP 3 673 419 B1**

(12)

**EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:  
**07.08.2024 Bulletin 2024/32**

(51) International Patent Classification (IPC):  
**G06N 3/098 (2023.01) G06N 3/0985 (2023.01)**  
**G06N 3/086 (2023.01) G06N 3/045 (2023.01)**

(21) Application number: **18807623.6**

(52) Cooperative Patent Classification (CPC):  
**G06N 3/098; G06N 3/086; G06N 3/0985;**  
**G06N 3/045**

(22) Date of filing: **22.11.2018**

(86) International application number:  
**PCT/EP2018/082162**

(87) International publication number:  
**WO 2019/101836 (31.05.2019 Gazette 2019/22)**

---

(54) **POPULATION BASED TRAINING OF NEURAL NETWORKS**

POPULATIONSBASIERTES TRAINING VON NEURONALEN NETZEN

APPRENTISSAGE BASÉ SUR UNE POPULATION DE RÉSEAUX NEURONAUX

(84) Designated Contracting States:  
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR**

(74) Representative: **Marks & Clerk GST**  
**1 New York Street**  
**Manchester M1 4HD (GB)**

(30) Priority: **22.11.2017 US 201762590177 P**

(56) References cited:  
**EP-A2- 1 557 788 CN-A- 107 229 972**  
**CN-A- 107 301 459 US-A1- 2018 260 714**  
**US-B1- 6 523 016**

(43) Date of publication of application:  
**01.07.2020 Bulletin 2020/27**

(60) Divisional application:  
**24172554.8 / 4 383 136**

- **CHRISANTHA FERNANDO ET AL: "Convolution by Evolution: Differentiable Pattern Producing Networks", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 8 June 2016 (2016-06-08), XP080706752**
- **MAX JADERBERG ET AL: "Population Based Training of Neural Networks", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 27 November 2017 (2017-11-27), XP080840551**

(73) Proprietor: **Deepmind Technologies Limited**  
**London EC4A 3TW (GB)**

- (72) Inventors:
- **JADERBERG, Maxwell Elliot**  
**London N1C 4AG (GB)**
  - **CZARNECKI, Wojciech**  
**London N1C 4AG (GB)**
  - **GREEN, Timothy Frederick Goldie**  
**London N1C 4AG (GB)**
  - **DALIBARD, Valentin Clement**  
**London N1C 4AG (GB)**

**EP 3 673 419 B1**

---

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

**Description**

## CROSS-REFERENCE TO RELATED APPLICATIONS

5 **[0001]** This application claims priority to Provisional Application No. 62/590,177, filed on November 22, 2017.

## BACKGROUND

**[0002]** This specification relates to training neural networks.

10 **[0003]** Neural networks are machine learning models that employ one or more layers of nonlinear units to predict an output for a received input. Some neural networks include one or more hidden layers in addition to an output layer. The output of each hidden layer is used as input to the next layer in the network (i.e., the next hidden layer or the output layer). Each layer of the network generates an output from a received input in accordance with current values of a respective set of parameters. Fernando et. al., "Convolution by Evolution: Differentiable Pattern Producing Networks," arXiv:1606.02580v1, 8 June 2016, available at <https://arxiv.org/abs/1606.02580>, discloses a differentiable version of the Compositional Pattern Producing Network, called the DPPN. Unlike a standard CPPN, the topology of a DPPN is evolved but the weights are learned. A Lamarckian algorithm, that combines evolution and learning, produces DPPNs to reconstruct an image.

## 20 SUMMARY

**[0004]** The invention is defined by the independent claims. Preferred embodiments are defined by the dependent claims.

25 **[0005]** Generally, this specification describes a distributed system comprising a plurality of computing units that trains a neural network having a plurality of network parameters to perform a particular neural network task. The system trains the neural network to determine trained values of the network parameters using an iterative training process that has a plurality of hyperparameters.

30 **[0006]** During the training of the neural network, the system maintains a plurality of candidate neural networks and, for each of the candidate neural networks, data specifying: (i) respective values of the network parameters for the candidate neural network, (ii) respective values of the hyperparameters for the candidate neural network, and (iii) a quality measure that measures a performance of the candidate neural network on the particular neural network task. The candidate neural networks that are maintained are collectively referred to as the population in this specification.

**[0007]** In order to train the neural network, the system repeatedly performs a set of training operations for each of the plurality of candidate neural networks.

35 **[0008]** As part of performing the training operations, the system trains the candidate neural network using the iterative training process and in accordance with the maintained values of the hyperparameters for the candidate neural network until termination criteria are satisfied to determine updated values of the network parameters for the candidate neural network from the maintained values of the network parameters for the candidate neural network.

40 **[0009]** The system then determines an updated quality measure for the candidate neural network in accordance with the updated values of the network parameters for the candidate neural network and the maintained values of the hyperparameters for the candidate neural network and determining, based at least on the maintained quality measures for the candidate neural networks and the updated quality measure for the candidate neural network, new values of the hyperparameters and the network parameters for the candidate neural network.

45 **[0010]** The system determines a new quality measure for the candidate neural network in accordance with the new values of the network parameters for the candidate neural network and the new values of the hyperparameters for the candidate neural network and updates the maintained data for the candidate neural network to specify the new values of the hyperparameters, the new values of the network parameters, and the new quality measure.

50 **[0011]** After repeatedly performing the set of training operations, the system selects the trained values of the network parameters from the parameter values in the maintained data based on the maintained quality measures for the candidate neural networks after the training operations have repeatedly been performed.

**[0012]** Generally, the network parameters are values that impact the operations performed by the neural network and that are adjusted as part of the iterative training process. For example, the network parameters can include values of weight matrices and, in some cases, bias vectors, of the layers of the neural network.

55 **[0013]** The hyperparameters are values that are not modified by the iterative training process. The hyperparameters can include values that impact how the values of the network parameters are updated by the training process e.g., the learning rate or other update rule that defines how the gradients determine at the current training iteration are used to update the network parameter values, objective function values, e.g., entropy cost, weights assigned to various terms of the objective function, and so on.

**[0014]** The neural network can be configured to receive any kind of digital data input and to generate any kind of score, classification, or regression output based on the input.

**[0015]** For example, if the inputs to the neural network are images or features that have been extracted from images, the output generated by the neural network for a given image may be scores for each of a set of object categories, with each score representing an estimated likelihood that the image contains an image of an object belonging to the category.

**[0016]** As another example, if the input to the neural network is data characterizing the state of an environment being interacted with by an agent, e.g., a robot or other mechanical agent, the output generated by the neural network can be a policy output that defines a control input for the agent. For example, the output can include or define a respective probability for each action in a set of possible actions to be performed by the agent or a respective Q value, i.e., a return estimate, for each action in the set of possible actions. As another example, the output can identify a control input in a continuous space of control inputs.

**[0017]** As another example, if the inputs to the neural network are Internet resources (e.g., web pages), documents, or portions of documents or features extracted from Internet resources, documents, or portions of documents, the output generated by the neural network for a given Internet resource, document, or portion of a document may be a score for each of a set of topics, with each score representing an estimated likelihood that the Internet resource, document, or document portion is about the topic.

**[0018]** As another example, if the inputs to the neural network are features of an impression context for a particular advertisement, the output generated by the neural network may be a score that represents an estimated likelihood that the particular advertisement will be clicked on.

**[0019]** As another example, if the inputs to the neural network are features of a personalized recommendation for a user, e.g., features characterizing the context for the recommendation, e.g., features characterizing previous actions taken by the user, the output generated by the neural network may be a score for each of a set of content items, with each score representing an estimated likelihood that the user will respond favorably to being recommended the content item.

**[0020]** As another example, if the input to the neural network is a sequence of text in one language, the output generated by the neural network may be a score for each of a set of pieces of text in another language, with each score representing an estimated likelihood that the piece of text in the other language is a proper translation of the input text into the other language.

**[0021]** As another example, if the input to the neural network is a sequence representing a spoken utterance, the output generated by the neural network may be a score for each of a set of pieces of text, each score representing an estimated likelihood that the piece of text is the correct transcript for the utterance.

**[0022]** According to an aspect, there is provided a method according to claim 1.

**[0023]** The method may further comprise providing the trained values of the network parameters for use in processing new inputs to the neural network.

**[0024]** Selecting the trained values of the network parameters from the parameter values in the maintained data based on the maintained quality measures for the candidate neural networks may comprise: selecting the maintained parameter values of the candidate neural network having a best maintained quality measure of any of the candidate neural networks after the training operations have repeatedly been performed.

**[0025]** Selecting the trained values of the network parameters from the parameter values in the maintained data based on the maintained quality measures for the candidate neural networks may comprise: determining if the maintained quality measure for the respective candidate neural network in the plurality of candidate neural networks is below a first pre-determined percentage of the candidate neural networks and in response, sampling the maintained parameter values of a candidate neural network having a maintained quality measure above a second pre-determined percentage of the candidate neural networks, after the training operations have repeatedly been performed.

**[0026]** Repeatedly performing the training operations comprises repeatedly performing the training operations in parallel for each candidate neural network.

**[0027]** Repeatedly performing the training operations comprises repeatedly performing the training operations for each candidate neural network asynchronously from performing the training operations for each other candidate neural network. That is, the training operations performed for a candidate neural network is performed independently of the training operations of other candidate neural networks and does not require centralized control of the training operations.

**[0028]** Determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network may comprise: determining whether the updated quality measure is better than every maintained quality measure; and in response to determining that the updated quality measure is not better than every maintained quality measure, setting the new values of the network parameters to the maintained values of the network parameters for the candidate neural network having the best maintained quality measure.

**[0029]** Optionally, in response to determining that the updated quality measure is better than every maintained quality measure, setting the new values of the network parameters to the updated values of the network parameters.

**[0030]** Determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network may comprise: setting the new values of the hyperparameters based upon the maintained values of the hyperparameters for the candidate neural network having the best maintained quality measure.

**[0031]** Setting the values of the hyperparameters may comprise randomly permuting one or more of the maintained values of the hyperparameters of the candidate neural network having the best maintained quality measure.

**[0032]** Determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network may comprise: sampling the new values of the hyperparameters from a prior distribution over possible values for the hyperparameters.

**[0033]** Determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network may comprise: sampling the new values of the network parameters from a prior distribution over possible values for the network parameters.

**[0034]** Determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network may comprise: determining whether the updated quality measure is better than every maintained quality measure; and in response to determining that the updated quality measure is not better than every maintained quality measure, setting the new values of the hyperparameters to the maintained values of the hyperparameters for the candidate neural network having the best maintained quality measure.

**[0035]** Training the candidate neural network using the iterative training process and in accordance with the maintained values of the hyperparameters for the candidate neural network until termination criteria are satisfied may comprise training the candidate neural network until a threshold number of iterations of the iterative training process have been performed.

**[0036]** As will be appreciated, the method of training a neural network described above may be implemented by the system for training a neural network.

**[0037]** The subject matter described in this specification can be implemented in particular implementations so as to realize one or more of the following advantages. By training the neural network in a manner that optimizes parameters and hyperparameters jointly, the system and method can train the neural network to generate network outputs that are more accurate than those generated by networks trained using conventional techniques. Additionally, the system and method can train the neural network more quickly, i.e., in terms of wall clock time, than other approaches.

**[0038]** Additionally, compared to other approaches that first optimize hyperparameters and then train the neural network to optimize the network parameters, the approach described in this specification uses fewer computational resources, i.e., less processing power and processing time, because the hyperparameters and network parameters are optimized jointly. In particular, many conventional approaches first search for acceptable values of the hyperparameters and then begin the training process once acceptable hyperparameter values have been found. This search for acceptable hyperparameter values can be computationally expensive and involve evaluating many possible combinations of hyperparameter values. By effectively optimizing the hyperparameters jointly with the network parameters as described in this specification, this computationally expensive stage of neural network training can be eliminated.

**[0039]** Additionally, the training is performed by a distributed system and the described approach only requires values of parameters, hyperparameters, and quality measures to be communicated between candidates, reducing the amount of data that needs to be communicated over the network during the training of the neural network. Moreover, the training operations are performed asynchronously and in a decentralized manner for each candidate neural network, making it so that the described training technique requires minimal overhead and infrastructure to be used to effectively train the neural network. As such, the training method is specifically adapted to take advantage of parallel processing systems but does not require centralized control of operations or for large amounts of data to be communicated between processing units of the parallel processing system. The method enables neural network training to be carried out more efficiently and to produce a better trained neural network.

**[0040]** As will be apparent, the system and method for training a neural network is universally applicable to any type of neural network for any type of technical task that neural networks may be applied to. For example, the system and method may be used to train a neural network for processing image data and video data, for example to recognize objects or persons in images and video. The network may be used for processing of audio signals, such as speech signals for performing speech recognition, speaker recognition and authentication and spoken language translation. The network may be used to control agent interacting in an environment, for example, a robot operating in a warehouse or an autonomous vehicle operating in the real world. The neural network may be used to perform data encoding or compression.

**[0041]** The details of one or more implementations of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0042]**

- 5           FIG. 1 shows an example population based neural network training system.  
           FIG. 2 is a flow chart of an example process for updating a population repository for a candidate neural network.  
           FIG. 3 shows an example candidate neural network before and after updating the population repository for the candidate neural network.

## 10   DETAILED DESCRIPTION

**[0043]** FIG. 1 shows an example population based neural network training system 100 ("the system 100").

**[0044]** The population based neural network training system 100 is a system that receives, (e.g., from a user of the system) training data 105 for training a neural network to perform a machine learning task. Training data 105 includes multiple training examples and a respective target output for each training example. The target output for a given training example is the output that should be generated by a trained neural network by processing the given training example.

**[0045]** The system 100 can receive the training data 105 in any of a variety of ways. For example, the system 100 can receive training data as an upload from a remote user of the system over a data communication network (e.g., using an application programming interface (API) made available by the system 100). As another example, the system 100 can receive an input from a user specifying which data that is already maintained by the system 100 should be used as the training data 105.

**[0046]** The system 100 generates data 150 specifying a trained neural network using the training data 105. The data 150 specifies the architecture of the trained neural network and the trained values of the parameters of the trained neural network. The parameters of the neural network will be referred to in this specification as "network parameters."

**[0047]** In some implementations, once the neural network has been trained, the system 100 provides data specifying the trained neural network 150 for use in processing new network inputs. That is, the system 100 can output (e.g., by outputting to a user device or by storing in memory accessible to the system) the trained values of the network parameters of the trained neural network 150 for later use in processing inputs using the trained neural network 150. Alternatively or in addition to outputting the trained neural network data 150, the system 100 can instantiate an instance of the neural network having the trained values of the network parameters, receive inputs to be processed, (e.g., through an API offered by the system 100) use the trained neural network 150 to process the received inputs to generate outputs, and then provide the generated outputs in response to the received inputs.

**[0048]** The system 100 includes a population repository 110 storing a plurality of candidate neural networks 120A-N (referred to in this specification as the "population"). The population repository 110 is implemented as one or more logical storage devices in one or more physical locations or as logical storage space allocated in one or more storage devices in one or more physical locations. At any given time during training, the repository 110 stores data specifying the current population of the candidate neural networks 120A-N. In some implementations, the population size (e.g., the number of candidate neural networks trained by the system 100) is greater than ten candidate neural networks. In some implementations, the population size is between twenty and eighty candidate neural networks.

**[0049]** The population repository 110 stores, for each candidate neural network 120A-N in the current population, a set of maintained values that defines the respective candidate neural network. The set of maintained values includes network parameters, hyperparameters, and a quality measure for each candidate neural network 120A-N (e.g., for candidate neural network A 120A, the set of maintained values includes network parameters A 125A, hyperparameters A 130A, and quality measure A 135A). During training, the network parameters, the hyperparameters, and the quality measure for a candidate neural network are updated in accordance with training operations, including an iterative training process (discussed below).

**[0050]** Generally, the network parameters are values that impact the operations performed by the candidate neural network and are adjusted as part of the iterative training process. For example, the network parameters can include values of weight matrices and, in some cases, bias vectors, of the layers of the candidate neural network. As another example, the network parameters can include values of kernels of convolutional layers in the candidate neural network.

**[0051]** The hyperparameters for a candidate neural network are values that are not modified by the iterative training process. The hyperparameters can include values that impact how the values of the network parameters are updated by training (e.g., learning rate, objective function values, entropy cost, auxiliary loss weights, or weights assigned to various terms of the objective function).

**[0052]** The quality measure for a candidate neural network is a measure of the performance of the candidate neural network on the particular machine learning task. Generally, the function defining the quality measure for the candidate neural network is not required to be differentiable, nor is it required to be the same function used during the iterative training process for the candidate neural network (i.e., may be different from the objective function optimized during the

iterative training process).

**[0053]** The system 100 trains each candidate neural network 120A-N by repeatedly performing iterations of an iterative training process to determine updated network parameters for the respective candidate neural network. At certain points during the iterative training process the system 100 also updates the repository 110 by performing additional training operations (discussed below).

**[0054]** To begin the training process, the population based neural network training system 100 pre-populates the population repository 110 with a plurality of candidate neural networks 120A-N for performing the specified machine learning task. In some implementations, the system 100 randomly initializes network parameters 125A-N and hyperparameters 130A-N for each candidate neural network 120A-N.

**[0055]** Alternatively or in addition, the system 100 initializes hyperparameters 130A-N so that each candidate neural network 120A-N in the population is initialized with unique hyperparameters (i.e., for any hyperparameters  $i$  and hyperparameters  $j$  from the hyperparameters 130A-N,  $i$  is not equal to  $j$ ).

**[0056]** For example, the system 100 generates a list of sets of hyperparameters of incrementally increasing values (so that no two sets of hyperparameters are equal) and initializes each candidate neural network 120A-N with a set of hyperparameters in the list, one set of hyperparameters for each candidate neural network 120A-N.

**[0057]** As another example, the system 100 randomly initializes the hyperparameters for each candidate neural network 120A-N and then scans the candidate neural networks 120AN to determine if hyperparameters are repeated for any two candidate neural networks. If the system 100 determines two candidate neural networks were initialized with the same hyperparameters, the system generates new randomized hyperparameters for one of the candidate neural networks. The system 100 scans the candidate neural networks 120A-N and generates new randomized hyperparameters until no two candidate neural networks have the same hyperparameters.

**[0058]** Each candidate neural network 120A-N is an architecture that receives inputs that conform to the machine learning task (i.e., inputs that have the format and structure of the training examples in the training data 105) and generates outputs that conform to the machine

learning task (i.e., outputs that have the format and structure of the target outputs in the training data 105).

**[0059]** For some machine learning tasks, each candidate neural network 120A-N needs to be trained jointly with one or more other neural networks. For example, in a generative adversarial neural network machine learning task, the system 100 trains a candidate neural network with one other neural network (e.g., a candidate generator neural network and a candidate discriminator neural network). The system 100 then generates data specifying a pair of trained neural networks (e.g., a trained generator neural network and a trained discriminator neural network). For these machine learning tasks, the system 100 maintains for each candidate neural network 120A-N, the maintained values of the respective one or more other neural networks, in the population repository 110.

**[0060]** The system 100 executes training operations for each candidate neural network 120AN in parallel, asynchronously, and in a decentralized manner. In some implementations, each candidate neural network 120A-N is assigned a respective computing unit for executing population based training. The computing units are configured so that they can operate independently of each other. A computing unit may be, for example, a computer, a core within a computer having multiple cores, or other hardware or software within a computer capable of independently performing the computation required by the system 100 for executing the iterative training process and updating the repository 110 for each candidate neural network 120A-N.

**[0061]** For each of the candidate neural networks 120A-N in the current population, the system 100 executes an iterative training process for training each of the candidate neural networks 120A-N. Additional training operations are necessary to further optimize the network parameters and hyperparameters of a candidate neural network and are discussed with respect to FIG. 2, below.

**[0062]** The iterative training process optimizes the network parameters 125A-N for the population of candidate neural networks 120A-N. In some implementations, the iterative training process optimizes the network parameters 125A-N for the population of candidate neural networks 120A-N in an iterative manner by using a step function (e.g., stochastic gradient descent on some loss function). The step function receives as input the maintained hyperparameters and network parameters for the candidate neural network, and returns the candidate neural network with updated network parameters, in accordance with the maintained hyperparameters. For a single step, the step function can be represented mathematically as:  $\Theta' \leftarrow \text{step}(\Theta/h)$  for network parameters  $\Theta$ , hyperparameters  $h$  and updated network parameters  $\Theta'$ .

**[0063]** After criteria are satisfied for ending execution of the training operations, (i.e., the system 100 determines that training is over, e.g., based on some performance criteria) the system 100 selects an optimal candidate neural network from the candidate neural networks 120A-N. In particular, in some implementations, the system 100 selects the candidate neural network in the population that has the best quality measure. The system 100 can determine the candidate neural network in the population with the best quality measure by comparing each quality measure 135A-N for each candidate neural network 120A-N and, for example, selecting the candidate neural network with the highest quality measure.

**[0064]** The optimal candidate neural network selected is sometimes referred to in this specification as the "best" candidate neural network 120A-N in the population. A candidate neural network that has a higher quality measure than

another candidate neural network is considered "better" than the other candidate neural network.

**[0065]** In some implementations, the system 100 generates data specifying a trained neural network 150 from a plurality of optimized neural network candidates. For example, the system 100 compares the quality measures 135A-N for candidate neural networks 120A-N in the population, and selects one or more candidate neural networks 120A-N having a quality measure that meets or exceeds a certain pre-determined performance threshold. Then, the system 100 samples from the plurality of candidate neural networks selected and generates data 150 specifying a trained neural network having the network parameters of the sampled candidate neural network.

**[0066]** Repeating the iterative training process until a performance criteria is met for the candidate neural network is computationally expensive, due to the number of iterations (e.g., repeated calls to  $step(\theta|h)$ ) during training. In addition, incorrectly chosen hyperparameters for a candidate neural network can lead to un-optimal trained network parameters or potentially no trained network parameters at all (i.e., the iterative training process does not converge to a set of trained network parameters for the candidate neural network).

**[0067]** Therefore, the population based neural network training system 100 additionally performs meta-optimization, where the hyperparameters and network parameters for each candidate neural network 120A-N are additionally adapted according to the performance of the entire population.

**[0068]** Training operations on the candidate neural networks 120A-N by the system 100 include operations to update the population repository 110 with new network parameters, new hyperparameters, and a new quality measure for each candidate neural network 120A-N. By periodically updating the candidate neural networks 120A-N independent of the iterative training process, the candidate neural networks 120A-N benefit from performance of the population. In addition, updating the hyperparameters before training is over can be more efficient, at least because training does not have to finish before applying more optimal hyperparameters.

**[0069]** In particular, modification of network parameters and hyperparameters for a candidate neural network in accordance with implementations of the disclosed subject matter minimizes error propagated through the system 100 using incorrect or sub-optimal parameters, by periodically checking the performance of other candidate neural networks trained in parallel and copying parameters from more optimized candidate neural networks. By training the candidate neural network in a manner that optimizes network parameters and hyperparameters jointly, the system can generate the trained neural network 150 to produce network outputs that are more accurate than those generated by networks trained using conventional techniques. The combination of updating network parameters and hyperparameters jointly results in increased performance of the candidate neural networks as opposed to updating only the network parameters, or only the hyperparameters.

**[0070]** The population based neural network training system 100 benefits from local optimization by executing, asynchronously and in parallel, an iterative training process for each candidate neural network 120A-N in the population. In particular, because the system executes the training operations for each candidate neural network 120A-N in parallel, the system 100 trains and periodically updates parameters in one training run, resulting in a shorter wall clock time as opposed to training, updating hyperparameters, and then retraining the neural network for the new hyperparameters. Moreover, the training operations is performed asynchronously and in a decentralized manner for each candidate neural network 120A-N, making it so that the described training technique requires minimal overhead and infrastructure to be used to effectively train the neural network.

**[0071]** Additionally, in training a neural network specified to execute certain machine learning tasks, the best values for the hyperparameters for the neural network can shift significantly at different phases of learning (called a "non-stationary learning problem") in variable and difficult to predict ways. The system 100 schedules periodic hyperparameter shift during training (e.g., learning-rate annealing) automatically and without a massive design-space.

**[0072]** Implementations of the disclosed subject matter demonstrate that the described techniques and systems can be adapted to train neural networks configured to execute a variety of machine learning tasks, as an adaptive strategy to training. Additionally, when the training is distributed, the described approach only requires values of parameters, hyperparameters, and quality measures to be communicated between candidate neural networks 120A-N, reducing the amount of data that needs to be communicated over the network during the training of the neural network 150. In some implementations, the system 100 includes auxiliary losses in the loss function to regularize or otherwise bias the solutions found, or to shape learning dynamics to speed up training. These auxiliary losses can be included in the system 100 without spending a long time tuning weight schedules by hand. That is, weights between different terms in the loss function can be automatically adjusted during the meta-optimization process.

**[0073]** Implementations of the disclosed subject matter allow for model selection at a population level based on attributes not directly trained for by candidate neural networks in the population. In those implementations, the system incorporates the attributes not directly trained for in the quality measure for a candidate neural network. For example, for a machine learning task specifying an image classification task, specifying data for generating a trained neural network includes comparing each candidate neural network not only for accuracy in image classification, but also by evaluating how robust the candidate neural network is to adversarial attack. The system 100 trains candidate neural networks for improved fitness for image classification, but additional training operations include meta-optimization based on, for

example, a quality measure that also includes the candidate neural network's robustness to adversarial attack.

**[0074]** In some implementations, for a machine learning task specifying a machine translation task, specifying data for generating a trained neural network 150 includes comparing each candidate neural network 120A-N not only for maximum likelihood of accuracy of translation between the input sequence and output sequence, but also based on an evaluation of the translation according to its bilingual evaluation understory score (BLEU) score (discussed below).

**[0075]** The system 100 also reduces the complexity of training neural networks for machine learning tasks, where two or more neural networks are trained together (e.g., a generator neural network and a discriminator neural network in a generative adversarial neural network framework). In particular, hyperparameters can be set independent of each candidate neural network in the two or more neural networks trained in tandem, as opposed to maintaining one set of hyperparameters for both networks (which requires a more complicated structure for maintaining consistency across the neural networks during training, and is often inefficient).

**[0076]** FIG. 2 is a flow chart of an example process for updating the population repository for a candidate neural network. For convenience, the process 200 will be described as being performed by a system of one or more computers located in one or more locations. For example, a population based neural network system (e.g., the population based neural network training system 100 of FIG. 1) appropriately programmed in accordance with this specification, can perform the process 200.

**[0077]** The system repeatedly performs the process 200 for each candidate neural network in the population. In particular, the system repeatedly performs the process 200 for each candidate neural network asynchronously from performing the process for each other candidate neural network in the population.

**[0078]** The system trains the candidate neural network until termination criteria are satisfied (step 202). Specifically, the system trains the candidate neural network using the maintained values for the hyperparameters and network parameters of the candidate neural network, to iteratively generate updated network parameters for the candidate neural network. Termination criteria are one or more conditions set, that when met by a candidate neural network, cause the system to update the repository for the candidate neural network, with new network parameters, hyperparameters, and quality measure. An example of a termination criterion being met is when a candidate neural network has been training for a set period of time or a fixed number of iterations of the iterative training process. Another example of a termination criterion being met is when a candidate neural network falls below a certain performance threshold. In those cases, the system continues the process 200 to update the repository for the candidate neural network.

**[0079]** Next, the system determines an updated quality measure for the candidate neural network in accordance with the updated values of the network parameters for the candidate neural network and the maintained values of the hyperparameters for the candidate neural network (step 204).

**[0080]** Specifically, the system executes the iterative training process on the candidate neural network to obtain updated network parameters; however, the maintained quality measure of the candidate neural network is not updated during the iterative training process to reflect the potential performance increase of the candidate neural network as a result of the updated network parameters. Therefore, before the system can compare the performance of the candidate neural networks to determine new network parameters and hyperparameters of the candidate neural network, the system performs step 204 to determine the updated quality measure for the candidate neural network. Examples for how the system determines the updated quality measure for the candidate neural network are discussed below.

**[0081]** The system determines new values of the hyperparameters and network parameters for the candidate neural network (step 206). In some implementations, the system determines new values of the hyperparameters and network parameters for the candidate neural network based at least on the maintained quality measures for the population of candidate neural networks in the population repository and the updated quality measure of the candidate neural network determined at step 204.

**[0082]** For example, the system determines the "best" candidate neural network having the highest quality measure in the current population, and determines if the best candidate neural network is better than the candidate neural network (e.g., by comparing the quality measures of both candidate neural networks and determining which quality measure is higher). If the best candidate neural network is better than the candidate neural network, then the system determines new values of the network parameters for the candidate neural network by "exploiting" (i.e., taking) the network parameters of the best candidate neural network in the current population. If the candidate neural network is better than the best candidate neural network in the population, then the new network parameters for the candidate neural network are set to the updated network parameters of the candidate neural network (i.e., the system does not exploit the current population).

**[0083]** Alternatively or in addition, the system generates new values of the hyperparameters in the same way the system determines new values for the network parameters of the candidate neural network (i.e., the system "exploits" the hyperparameters of the best candidate neural network in the population and sets those hyperparameters as the hyperparameters of the candidate neural network).

**[0084]** Determining new network parameters for the candidate neural network by "exploiting" is represented mathematically as  $\Theta' \leftarrow \text{exploit}(\Theta)$ , for updated network parameters  $\Theta'$  of the candidate neural network and returns new

network parameters  $\Theta$ " for the candidate neural network. Other implementations of *exploit*(·) are discussed in accordance with examples below.

**[0085]** In some implementations, the system "explores" for new hyperparameters for the candidate neural network by determining new hyperparameters using a mathematical operation that takes the maintained hyperparameters of the candidate neural network as input. For example, the system determines new hyperparameters for the candidate neural network by randomly perturbing the hyperparameters of the candidate neural network with noise and setting the perturbed values as the new hyperparameters for the candidate neural network. Alternatively, the system determines new hyperparameters of the candidate neural network by sampling from a distribution over possible values for the hyperparameters. Alternatively, the system determines new hyperparameters for the candidate neural network by incrementing or decrementing the values of the hyperparameters of the candidate neural network and setting the incremented or decremented values as the new hyperparameters for the candidate neural network.

**[0086]** Alternatively or in addition, the system generates new values of the network parameters in the same way the system determines new values for the hyperparameters of the candidate neural network (i.e., the system "explores" for new network parameters by randomly perturbing the updated network parameters of the candidate neural network with noise and setting the perturbed values as the new network parameters for the candidate neural network).

**[0087]** In some implementations, the system first determines whether the candidate neural network is the best candidate neural network in the current population (e.g., by comparing the quality measure of the candidate neural network with the quality measures of the other candidate neural networks in the population). If the candidate neural network is the best candidate neural network, then the system does not "explore" for new hyperparameters for the candidate neural network.

**[0088]** In some implementations, the system determines new hyperparameters for the candidate neural network by "exploring" regardless of whether or not the candidate neural network is the best in the current population.

**[0089]** Determining new hyperparameters for the candidate neural network by "exploring" is represented mathematically as:  $h' \leftarrow \text{explore}(h)$  for hyperparameters  $h$  of the candidate neural network, and returns new hyperparameters  $h'$  for the candidate neural network. Other implementations of *explore*(·) are discussed in accordance with examples below.

**[0090]** In the exemplary process 200 for updating the population repository for the candidate neural network, the system updates the hyperparameters and network parameters jointly for the candidate neural network. In some implementations, the system may intermittently update the population repository with only new hyperparameters for the candidate neural network, or only new network parameters for the candidate neural network (e.g., for multiple iterations of the process 200 on the candidate neural network, the system may only update the repository with new hyperparameters or new network parameters for the candidate neural network after every other iteration of the process 200).

**[0091]** The system determines a new quality measure for the candidate neural network in accordance with the new values of the network parameters and hyperparameters for the candidate neural network (step 208). Specifically, the system takes the new values of the network parameters and hyperparameters for the candidate neural network determined in step 206 to determine the new quality measure for the candidate neural network.

**[0092]** The system updates the population repository with the new quality measure, network parameters, and hyperparameters for the candidate neural network (step 210). Specifically, the system updates the population repository with the new network parameters and hyperparameters for the candidate neural network determined in step 206, and the new quality measure for the candidate neural network determined in step 208.

**[0093]** After the system executes process 200 for the candidate neural network, the iterative training process continues. Specifically, the system trains the candidate neural network using the new hyperparameters and new network parameters of the candidate neural network, to iteratively generate updated network parameters for the candidate neural network.

**[0094]** The system will continue the iterative training process for the candidate neural network until either the termination criteria is satisfied (and the system repeats the process 200 for the candidate neural network) or performance criteria is satisfied to indicate to the system to stop training.

**[0095]** FIG. 3 shows an example candidate neural network before and after updating the population repository for the candidate neural network. For illustrative purposes, FIG. 3 shows a candidate neural network A 320A ("network A") and a candidate neural network 320B ("network B") before and after termination criteria for candidate neural network B 321. Network A 320A in this example is better than Network B 320B.

**[0096]** The system executes iterative training processes for network A 320A and network B 320B, in parallel (indicated by line connectors 350A and 350B, respectively). As the system 100 executes the iterative training processes for network A 320A and network B 320B, network parameters A 325A and network parameters B 325B are updated and replaced by updated network parameters A 340A and updated network parameters B 340B, respectively.

**[0097]** The system 100 determines an updated quality measure B 336B for the network B 320B with the values of the hyperparameters B 330B and the updated network parameters B 340B, (i.e., the system 100 executes step 204 of process 200 for network B 320B, represented as line connector 360 in FIG. 3). The system 100 determines new hyperparameters B 331B and new network parameters B 341B for the network B 320B (i.e., the system 100 executes step 206 of process 200 for network B 320B).

## EP 3 673 419 B1

**[0098]** The system 100 compares the quality measure A 335A and the quality measure B 335B, and sets new network parameters B 341B to be equal to the updated network parameters A 340A (represented as line connector 370), because in this example, network A 320A is better than network B 320B. The system 100 determines new hyperparameters B 341B for the network B 320B, represented by line connector 380.

**[0099]** The system 100 determines from the new network parameters B 341B and the new hyperparameters B 331B for the network B 320B, a new quality measure B 337B (i.e., the system 100 executes step 208 of process 200 for the network B 320B).

**[0100]** The system 100 updates the population repository 110 for the network B 320B to specify the new quality measure B 337B, new network parameters B 341B, and new hyperparameters B 331B (i.e., the system 100 executes step 210 of process 200 for the network B 320B). The system 100 resumes the iterative training process for network B 320B.

**[0101]** While the system 100 determined new values for the network B 320B and updated the repository 110, system 100 continued to train network 320A, in parallel. When training is over, the system 100 specifies data for trained neural network A 390A and trained neural network B 390B.

**[0102]** Examples of certain implementations of the described meta-optimization schemes follow. In the below examples,  $\cdot$  is a placeholder for a function  $f(\cdot)$  and represents zero or more arguments for the function  $f$ .

**[0103]** Training operations executed by a population based neural network training system for each candidate neural network are described with the following definitions:

*step*( $\cdot$ ): How the system in the exemplary implementation iterates one step in the iterative training process, for a candidate neural network.

*eval*( $\cdot$ ): How the system determines the quality measure for the candidate neural network.

*ready*( $\cdot$ ): The termination criteria for the system, represented as a function that returns a value indicating whether the termination criteria that cause the system to update the population repository are satisfied for the candidate neural network.

*exploit*( $\cdot$ ): How the system updates the population repository with new values for the network parameters of the candidate neural network.

*explore*( $\cdot$ ): How the system updates the population repository with new values for the hyperparameters of the candidate neural network.

**[0104]** In one example, a reinforcement learning task is specified as the machine learning task for the system. Each candidate neural network receives inputs and generates outputs that conform to a deep reinforcement learning task. In reinforcement learning, the aim is to find a policy to maximize expected episodic return within an environment. An exemplary set of training operations for population based neural network training for a deep reinforcement machine learning task includes:

*step*( $\cdot$ ): The system executes gradient descent on the values of the weights in the network parameters for the candidate neural network to determine updated network parameters (e.g., RMSProp).

*eval*( $\cdot$ ): The system updates the quality measure of a candidate neural network based on the mean value of a pre-determined number of previous episodic rewards (e.g., 10 episodic rewards). Specifically, the quality measure of the candidate neural network is the mean value of the pre-determined number of previous episodic rewards. The candidate neural network with the highest mean episodic reward has the highest quality measure and is considered the "best" in terms of measured fitness.

*ready*( $\cdot$ ): Termination criteria are satisfied after a fixed number of iterations (e.g.,  $1 \times 10^6$  to  $10 \times 10^6$  iterations of *step*( $\cdot$ )).

*exploit*( $\cdot$ ): The system uniformly samples another candidate neural network in the population and compares the last 10 episode rewards using Welch's  $t$ -test (also referred to as Welch's unequal variances  $t$ -test). If the sampled candidate neural network has a higher quality measure and satisfies the  $t$ -test, the system 100 sets the new network parameters and new hyperparameters of the candidate neural network to the values of the sampled candidate neural network.

**[0105]** Alternatively, the system ranks all candidate neural networks in the population by mean episode reward (i.e., rank each candidate neural network by ordering each quality measure from best to worst). If the candidate neural network is ranked below a certain threshold (e.g., bottom 20% of all candidate neural networks), the system 100 samples a candidate neural network ranked above a certain threshold (e.g., top 20% of all candidate neural networks) and sets the new network parameters and new hyperparameters of the candidate neural network to the values of the sampled candidate neural network. If the candidate neural network is ranked better than the bottom 20%, then the system does not exploit the population for new network parameters and new hyperparameters for the candidate neural network.

**[0106]** *explore*(·): The system randomly perturbs the hyperparameters of the candidate neural network by a pre-determined factor (e.g., by a factor of 1.2 or .8) and sets the perturbed hyperparameters as the new hyperparameters for the candidate neural network.

**[0107]** Alternatively, the system resamples the hyperparameters from the original prior distribution of hyperparameters for the candidate neural network, according to some probability, and sets the resampled hyperparameters as the new hyperparameters of the candidate neural network.

**[0108]** In another example, a supervised learning task is specified as the machine learning task for the system. Each candidate neural network receives inputs and generates outputs that conform to a supervised machine learning task. One example of a supervised learning task is machine translation from one language to another language. In particular, the task is to encode a source sequence of words in one language, and output a sequence in a different target language. An exemplary set of training operations for population based neural network training of a machine translation learning task includes:

*step*(·): The system executes gradient descent on the values of the weights in the network parameters for a candidate neural network, with adaptive moment estimation, (e.g., ADAM) to determine updated network parameters of the candidate neural network.

*eval*(·): The system updates the quality measure of the candidate neural network based on evaluation of the candidate neural network's bilingual evaluation understudy score (BLEU score). Specifically, the quality measure for a candidate neural network is its BLEU score. The BLEU score is an evaluation of a generated sequence compared with a reference sequence. A BLEU score of 1.0 is a perfect match between the generated sequence and the reference sequence, and a BLEU score of 0.0 is a perfect mis-match between the generated sequence and the reference sequence. The candidate neural network with the highest mean BLEU score has the highest quality measure and is considered the "best" in terms of measured fitness. Because the BLEU score is not differentiable, the BLEU score cannot directly be used as the loss function optimized during the iterative training process.

*ready*(·): Termination criteria are satisfied after a fixed number of iterations (e.g.,  $2 \times 10^8$  iterations of *step*(·)).

*exploit*(·): The system uniformly samples another candidate neural network in the population and compares the last 10 BLEU scores using Welch's *t*-test. If the sampled candidate neural network has a higher quality measure and satisfies the *t*-test, set the new network parameters and new hyperparameters of the candidate neural network to the values of the sampled candidate neural network.

*explore*(·): The system randomly perturbs the hyperparameters of the candidate neural network by a pre-determined factor (e.g., by a factor of 1.2 or .8) and sets the perturbed hyperparameters as the new hyperparameters for the candidate neural network.

**[0109]** Alternatively, the system resamples the hyperparameters from the original distribution of hyperparameters for the candidate neural network, according to some probability, and sets the resampled hyperparameters as the new hyperparameters of the candidate neural network.

**[0110]** In some implementations, a generative adversarial network (GAN) task is specified as the machine learning task for the population based neural network training system. Each candidate neural network receives inputs and generates outputs that conform to a GAN task. In a GAN task, two computing modules, represented as separate neural networks are defined: a generator neural network and a discriminator neural network. The discriminator takes as input samples from the generator and a real data distribution, and is trained to predict whether inputs are real or generated. For example, the generator maps simulated images to real-world images with the objective of maximally fooling the discriminator into classifying or scoring its images as real. An exemplary set of training operations for a population of pairs of candidate generator and discriminator neural networks of a GAN task includes:

*step*(·): For the candidate discriminator neural network, the system executes gradient descent a pre-determined number of iterations (e.g., the system executes gradient descent 5 times on the candidate discriminator neural network). The system updates the candidate generator neural network once using the ADAM optimizer. The system trains using a Wasserstein GAN gradient penalty (WGAN-GP) (Gulrajani, et. al) objective where the candidate discriminator neural network estimates the Wasserstein distance between the real and generated data distributions, with the Lipschitz constraint enforced by regularizing its input gradient to have a unit norm.

*eval*(·): The candidate discriminator neural network and the candidate generator neural network are evaluated by a variant of the Inception score proposed by Salimans et al., computed from the outputs of a pre-trained classifier. The candidate discriminator neural network and the candidate generator neural network with the better Inception score has the highest quality measure and is considered the "best" in terms of measured fitness.

*ready*(·): Termination criteria are satisfied after a fixed number of iterations (e.g.,  $1 \times 10^4$  iterations of *step*(·)).

*exploit*(·): The system ranks each pair of candidate generator neural network and discriminator neural networks by ordering each respective pair of quality measures from best to worst. If the pair of candidate neural networks is

ranked below a certain threshold (e.g., bottom 20% of all pairs of candidate neural networks), the system samples another pair of candidate neural networks ranked above a certain threshold (e.g., top 20% of all pairs candidate neural networks) and sets the new network parameters and new hyperparameters of the generator candidate neural network to the values of the network parameters and hyperparameters of the sampled generator candidate neural network and sets the new network parameters and new hyperparameters of the discriminator candidate neural network to the values of the network parameters and hyperparameters of the sampled discriminator candidate neural network.

**[0111]** Alternatively, for each pair of candidate neural networks, the system randomly samples a competing pair of candidate neural networks in the population. If the competing pair of candidate neural networks has a higher quality measure than the pair of candidate neural networks, the system sets the new network parameters and new hyperparameters of the generator and discriminator candidate neural networks to the values of the network parameters and hyperparameters of the competing generator and discriminator candidate neural networks, respectively.

**[0112]** *explore*( $\cdot$ ): The system randomly perturbs the hyperparameters of the generator and the discriminator candidate neural network by a pre-determined factor (e.g., by a factor of 2.0 or .5) and sets the perturbed hyperparameters for the candidate generator neural network and discriminator neural network as the new hyperparameters for the candidate generator and discriminator neural networks, respectively.

**[0113]** TABLE 1 illustrates an example of a procedure for executing training operations for a population of candidate neural networks:

Algorithm 1 Population Based Training (PBT)	
1: procedure TRAIN( $\mathcal{P}$ )	$\triangleright$ initial population $\mathcal{P}$
2: for $(\theta, h, p, t) \in \mathcal{P}$ (asynchronously in parallel) do	
3: while not end of training do	
4: $\theta \leftarrow \text{step}(\theta h)$	$\triangleright$ one step of optimisation using hyperparameters $h$
5: $p \leftarrow \text{eval}(\theta)$	$\triangleright$ current model evaluation
6: if $\text{ready}(p, t, \mathcal{P})$ then	
7: $h', \theta' \leftarrow \text{exploit}(h, \theta, p, \mathcal{P})$	$\triangleright$ use the rest of population to find better solution
8: if $\theta \neq \theta'$ then	
9: $h, \theta \leftarrow \text{explore}(h', \theta', \mathcal{P})$	$\triangleright$ produce new optimisation hyperparameters $h$
10: $p \leftarrow \text{eval}(\theta)$	$\triangleright$ new model evaluation
11: end if	
12: end if	
13: update $\mathcal{P}$ with new $(\theta, h, p, t + 1)$	$\triangleright$ update population
14: end while	
15: end for	
16: return $\theta$ with the highest $p$ in $\mathcal{P}$	
17: end procedure	

TABLE 1

**[0114]** *Train*( $\cdot$ ) is an exemplary function a population based neural network training system (e.g., the system 100 of FIG. 1) executes on a population of candidate neural networks. The system executes *Train*( $\mathcal{P}$ ), where  $\mathcal{P}$  is the population of candidate neural networks, and the population repository stores values for the network parameters  $\theta$ , hyperparameters  $h$ , and quality measure  $p$  for each candidate neural network. The system tracks a counter  $t$  that is incremented each time the population repository is updated.

**[0115]** The system updates the network parameters for a candidate neural network in the population (shown as  $\theta \leftarrow \text{step}(\theta|h)$  in line 4 in TABLE 1). After each iteration of the iterative training process, the system updates the quality measure for the candidate neural network (shown as  $p \leftarrow \text{eval}(\theta)$  in line 5 in TABLE 1).

**[0116]** Next, the system checks if the termination criteria are satisfied (shown as a conditional statement in line 6 in TABLE 1 executing *ready*( $p, t, \mathcal{P}$ )). If termination criteria are satisfied, the system 100 executes:  $h', \theta' \leftarrow \text{exploit}(h, \theta, p, \mathcal{P})$  for new hyperparameters  $h'$  and network parameters  $\theta'$  (shown in line 7 in TABLE 1). If the new network parameters and maintained network parameters of the candidate neural network are not equal, then the system executes  $h, \theta \leftarrow \text{explore}(h', \theta', \mathcal{P})$  to determine new hyperparameters  $h$  and new network parameters  $\theta$  for the candidate neural network (shown in line 8-9 in TABLE 1). The system updates the quality measure for the candidate neural network again (shown in line 10 in TABLE 1) and the population repository is updated to specify the new values for the hyperparameters, network parameters, and quality measure for the candidate neural network (shown in line 13 in TABLE 1). When training is over, the system generates data specifying a trained neural network by selecting the candidate neural network from the population with the highest quality measure (shown in line 16 in TABLE 1).

**[0117]** This specification uses the term "configured" in connection with systems and computer program components.

For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions.

**[0118]** Implementations of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Implementations of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non transitory storage medium for execution by, or to control the operation of, data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. Alternatively or in addition, the program instructions can be encoded on an artificially generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus.

**[0119]** The term "data processing apparatus" refers to data processing hardware and encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can also be, or further include, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit). The apparatus can optionally include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

**[0120]** A computer program, which may also be referred to or described as a program, software, a software application, an app, a module, a software module, a script, or code, can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages; and it can be deployed in any form, including as a stand alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a data communication network.

**[0121]** In this specification, the term "database" is used broadly to refer to any collection of data: the data does not need to be structured in any particular way, or structured at all, and it can be stored on storage devices in one or more locations. Thus, for example, the index database can include multiple collections of data, each of which may be organized and accessed differently.

**[0122]** Similarly, in this specification the term "engine" is used broadly to refer to a software-based system, subsystem, or process that is programmed to perform one or more specific functions. Generally, an engine will be implemented as one or more software modules or components, installed on one or more computers in one or more locations. In some cases, one or more computers will be dedicated to a particular engine; in other cases, multiple engines can be installed and running on the same computer or computers.

**[0123]** The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by special purpose logic circuitry, e.g., an FPGA or an ASIC, or by a combination of special purpose logic circuitry and one or more programmed computers.

**[0124]** Computers suitable for the execution of a computer program can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. The central processing unit and the memory can be supplemented by, or incorporated in, special purpose logic circuitry. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

**[0125]** Computer readable media suitable for storing computer program instructions and data include all forms of non volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM,

EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks.

**[0126]** To provide for interaction with a user, implementations of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's device in response to requests received from the web browser. Also, a computer can interact with a user by sending text messages or other forms of message to a personal device, e.g., a smartphone that is running a messaging application, and receiving responsive messages from the user in return.

**[0127]** Data processing apparatus for implementing machine learning models can also include, for example, special-purpose hardware accelerator units for processing common and compute-intensive parts of machine learning training or production, i.e., inference, workloads.

**[0128]** Machine learning models can be implemented and deployed using a machine learning framework, e.g., a TensorFlow framework, a Microsoft Cognitive Toolkit framework, an Apache Singa framework, or an Apache MXNet framework.

**[0129]** Implementations of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface, a web browser, or an app through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (LAN) and a wide area network (WAN), e.g., the Internet.

**[0130]** The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some implementations, a server transmits data, e.g., an HTML page, to a user device, e.g., for purposes of displaying data to and receiving user input from a user interacting with the device, which acts as a client. Data generated at the user device, e.g., a result of the user interaction, can be received at the server from the device.

**[0131]** While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or on the scope of what may be claimed, but rather as descriptions of features that may be specific to particular implementations of particular inventions.

**[0132]** Similarly, while operations are depicted in the drawings and recited in the claims in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

**[0133]** Particular implementations of the subject matter have been described. Other implementations are within the scope of the following claims. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In some cases, multitasking and parallel processing may be advantageous.

## Claims

1. A method of training a neural network having a plurality of network parameters to perform a particular neural network task and to determine trained values of the network parameters using an iterative training process having a plurality of hyperparameters, the method implemented by a distributed system comprising a plurality of computing units, the method comprising:

maintaining a plurality of candidate neural networks (120A, 120N) and, for each of the candidate neural networks, data specifying: (i) respective values of the network parameters (125A, 125N) for the candidate neural network, (ii) respective values of the hyperparameters (130A, 130N) for the candidate neural network, and (iii) a quality

measure (135A, 135N) that measures a performance of the candidate neural network on the particular neural network task;

for each candidate neural network of the plurality of candidate neural networks (120A, 120N), assigning a respective computing unit and repeatedly performing, by the assigned computing unit, the following training operations, the training operations performed in parallel for each candidate neural network and asynchronously from performing the training operations for each other candidate neural network:

training (202) the candidate neural network using the iterative training process and in accordance with the maintained values of the hyperparameters for the candidate neural network until termination criteria are satisfied to determine updated values of the network parameters for the candidate neural network from the maintained values of the network parameters for the candidate neural network,

determining (204) an updated quality measure for the candidate neural network in accordance with the updated values of the network parameters for the candidate neural network and the maintained values of the hyperparameters for the candidate neural network,

determining (206), based at least on the maintained quality measures for the candidate neural networks and the updated quality measure for the candidate neural network, new values of the hyperparameters and the network parameters for the candidate neural network,

determining (208) a new quality measure for the candidate neural network in accordance with the new values of the network parameters for the candidate neural network and the new values of the hyperparameters for the candidate neural network, and

updating (210) the maintained data for the candidate neural network to specify the new values of the hyperparameters, the new values of the network parameters, and the new quality measure; and

selecting the trained values of the network parameters from the parameter values in the maintained data based on the maintained quality measures for the candidate neural networks after the training operations have repeatedly been performed.

**2.** The method of claim 1, further comprising:

providing the trained values of the network parameters for use in processing new inputs to the neural network.

**3.** The method of any one of claims 1 or 2, wherein selecting the trained values of the network parameters from the parameter values in the maintained data based on the maintained quality measures for the candidate neural networks comprises:

selecting the maintained parameter values of the candidate neural network having a best maintained quality measure of any of the candidate neural networks after the training operations have repeatedly been performed.

**4.** The method of any of claims 1 or 2, wherein selecting the trained values of the network parameters from the parameter values in the maintained data based on the maintained quality measures for the candidate neural networks comprises:

determining if the maintained quality measure for the respective candidate neural network in the plurality of candidate neural networks is below a first pre-determined percentage of the candidate neural networks; and in response, sampling the maintained parameter values of a candidate neural network having a maintained quality measure above a second pre-determined percentage of the candidate neural networks, after the training operations have repeatedly been performed.

**5.** The method of any one of claims 1-4, wherein determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network comprises:

determining whether the updated quality measure is better than every maintained quality measure; and in response to determining that the updated quality measure is not better than every maintained quality measure, setting the new values of the network parameters to the maintained values of the network parameters for the candidate neural network having the best maintained quality measure.

**6.** The method of claim 5, further comprising:

in response to determining that the updated quality measure is better than every maintained quality measure, setting the new values of the network parameters to the updated values of the network parameters.

7. The method of any one of claims 1-6, wherein determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network comprises:  
5 setting the new values of the hyperparameters based upon the maintained values of the hyperparameters for the candidate neural network having the best maintained quality measure.
8. The method of claim 7, wherein setting the new values of the hyperparameters comprises randomly permuting one or more of the maintained values of the hyperparameters of the candidate neural network having the best maintained quality measure.  
10
9. The method of any one of claims 1-6, wherein determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network comprises:  
15 sampling the new values of the hyperparameters from a prior distribution over possible values for the hyperparameters; and/or  
sampling the new values of the network parameters from a prior distribution over possible values for the network parameters.
- 20 10. The method of any one of claims 1-4, wherein determining, based at least on the maintained quality measures for the candidate neural networks, new values of the hyperparameters and the network parameters for the candidate neural network comprises:  
25 determining whether the updated quality measure is better than every maintained quality measure; and  
in response to determining that the updated quality measure is not better than every maintained quality measure, setting the new values of the hyperparameters to the maintained values of the hyperparameters for the candidate neural network having the best maintained quality measure.
- 30 11. The method of any one of claims 1-10, wherein training the candidate neural network using the iterative training process and in accordance with the maintained values of the hyperparameters for the candidate neural network until termination criteria are satisfied comprises training the candidate neural network until a threshold number of iterations of the iterative training process have been performed.
- 35 12. A system comprising a plurality of computing units and one or more storage devices storing instructions that when executed by the system causes the plurality of computing units to perform the operations of the respective method of any one of claims 1-11.
- 40 13. One or more computer storage media storing instructions that when executed by a system comprising a plurality of computing units causes the plurality of computing units to perform the operations of the respective method of any one of claims 1-11.

## Patentansprüche

- 45 1. Verfahren zum Trainieren eines neuronalen Netzwerks mit einer Vielzahl von Netzwerkparametern, um eine bestimmte neuronale Netzwerkaufgabe auszuführen und um trainierte Werte der Netzwerkparameter unter Verwendung eines iterativen Trainingsprozesses mit einer Vielzahl von Hyperparametern zu bestimmen, wobei das Verfahren durch ein verteiltes System implementiert wird, das eine Vielzahl von Recheneinheiten umfasst, wobei das Verfahren Folgendes umfasst:  
50 Halten einer Vielzahl von neuronalen Kandidatennetzwerken (120A, 120N) und, für jedes der neuronalen Kandidatennetzwerke, von Daten, die Folgendes spezifizieren: (i) jeweilige Werte der Netzwerkparameter (125A, 125N) für das neuronale Kandidatennetzwerk, (ii) jeweilige Werte der Hyperparameter (130A, 130N) für das neuronale Kandidatennetzwerk, und (iii) ein Qualitätsmaß (135A, 135N), das eine Leistung des neuronalen Kandidatennetzwerks bei der bestimmten neuronalen Netzwerkaufgabe misst;  
55 Zuweisen, für jedes neuronale Kandidatennetzwerk der Vielzahl von neuronalen Kandidatennetzwerken (120A, 120N), einer jeweiligen Recheneinheit und wiederholtes Durchführen, durch die zugewiesene Recheneinheit, der folgenden Trainingsoperationen, wobei die Trainingsoperationen parallel für jedes neuronale Kandidaten-

netzwerk und asynchron zur Durchführung der Trainingsoperationen für jedes andere neuronale Kandidatennetzwerk durchgeführt werden:

5 Trainieren (202) des neuronalen Kandidatennetzwerks mittels des iterativen Trainingsprozesses und gemäß den gehaltenen Werten der Hyperparameter für das neuronale Kandidatennetzwerk, bis Terminierungskriterien erfüllt sind, um aktualisierte Werte der Netzwerkparameter für das neuronale Kandidatennetzwerk aus den gehaltenen Werten der Netzwerkparameter für das neuronale Kandidatennetzwerk zu bestimmen,  
10 Bestimmen (204) eines aktualisierten Qualitätsmaßes für das neuronale Kandidatennetzwerk gemäß den aktualisierten Werten der Netzwerkparameter für das neuronale Kandidatennetzwerk und den gehaltenen Werten der Hyperparameter für das neuronale Kandidatennetzwerk,  
Bestimmen (206), mindestens auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke und des aktualisierten Qualitätsmaßes für das neuronale Kandidatennetzwerk, neuer Werte der Hyperparameter und der Netzwerkparameter für das neuronale Kandidatennetzwerk,  
15 Bestimmen (208) eines neuen Qualitätsmaßes für das neuronale Kandidatennetzwerk gemäß den neuen Werten der Netzwerkparameter für das neuronale Kandidatennetzwerk und den neuen Werten der Hyperparameter für das neuronale Kandidatennetzwerk, und  
Aktualisieren (210) der gehaltenen Daten für das neuronale Kandidatennetzwerk, um die neuen Werte der Hyperparameter, die neuen Werte der Netzwerkparameter und das neue Qualitätsmaß zu spezifizieren; und  
20 Auswählen der trainierten Werte der Netzwerkparameter aus den Parameterwerten in den gehaltenen Daten auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke nach dem wiederholten Durchführen der Trainingsoperationen.

2. Verfahren nach Anspruch 1, das ferner Folgendes umfasst:  
25 Bereitstellen der trainierten Werte der Netzwerkparameter zur Verwendung bei der Verarbeitung neuer Eingaben in das neuronale Netzwerk.

3. Verfahren nach Anspruch 1 oder 2, wobei das Auswählen der trainierten Werte der Netzwerkparameter aus den Parameterwerten in den gehaltenen Daten auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetze Folgendes umfasst:  
30 Auswählen der gehaltenen Parameterwerte des neuronalen Kandidatennetzwerks mit einem besten gehaltenen Qualitätsmaß einer beliebigen der neuronalen Kandidatennetzwerke nach dem wiederholten Durchführen der Trainingsoperationen.

4. Verfahren nach Anspruch 1 oder 2, wobei das Auswählen der trainierten Werte der Netzwerkparameter aus den Parameterwerten in den gehaltenen Daten auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke Folgendes umfasst:

40 Feststellen, ob das gehaltene Qualitätsmaß für das jeweilige neuronale Kandidatennetzwerk in der Vielzahl von neuronalen Kandidatennetzwerken unter einem ersten vorbestimmten Prozentsatz der neuronalen Kandidatennetzwerke liegt; und  
Abtasten, als Reaktion darauf, der gehaltenen Parameterwerte eines neuronalen Kandidatennetzwerks mit einem gehaltenen Qualitätsmaß über einem zweiten vorbestimmten Prozentsatz der neuronalen Kandidatennetzwerke nach dem wiederholten Durchführen der Trainingsoperationen.

5. Verfahren nach einem der Ansprüche 1 bis 4, wobei das Bestimmen, mindestens auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke, neuer Werte der Hyperparameter und der Netzwerkparameter für das neuronale Kandidatennetzwerk Folgendes umfasst:

50 Feststellen, ob das aktualisierte Qualitätsmaß besser ist als jedes gehaltene Qualitätsmaß; und  
Setzen, als Reaktion auf die Feststellung, dass das aktualisierte Qualitätsmaß nicht besser als jedes gehaltene Qualitätsmaß ist, der neuen Werte der Netzwerkparameter auf die gehaltenen Werte der Netzwerkparameter für das neuronale Kandidatennetzwerk mit dem besten gehaltenen Qualitätsmaß.

6. Verfahren nach Anspruch 5, das ferner Folgendes umfasst:  
55 Setzen, als Reaktion auf die Feststellung, dass das aktualisierte Qualitätsmaß besser ist als jedes gehaltene Qualitätsmaß, der neuen Werte der Netzwerkparameter auf die aktualisierten Werte der Netzwerkparameter.

## EP 3 673 419 B1

7. Verfahren nach einem der Ansprüche 1 bis 6, wobei das Bestimmen, mindestens auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke, neuer Werte der Hyperparameter und der Netzwerkparameter für das neuronale Kandidatennetzwerk Folgendes umfasst:  
Setzen der neuen Werte der Hyperparameter auf der Basis der gehaltenen Werte der Hyperparameter für das neuronale Kandidatennetzwerk mit dem besten gehaltenen Qualitätsmaß.
8. Verfahren nach Anspruch 7, wobei das Setzen der neuen Werte der Hyperparameter das zufällige Permutieren eines oder mehrerer der gehaltenen Werte der Hyperparameter des neuronalen Kandidatennetzwerks mit dem besten gehaltenen Qualitätsmaß umfasst.
9. Verfahren nach einem der Ansprüche 1-6, wobei das Bestimmen, mindestens auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke, neuer Werte der Hyperparameter und der Netzwerkparameter für das neuronale Kandidatennetzwerk Folgendes umfasst:
- Abtasten der neuen Werte der Hyperparameter aus einer vorherigen Verteilung über mögliche Werte für die Hyperparameter; und/oder  
Abtasten der neuen Werte der Netzwerkparameter aus einer vorherigen Verteilung über mögliche Werte für die Netzwerkparameter.
10. Verfahren nach einem der Ansprüche 1 bis 4, wobei das Bestimmen, mindestens auf der Basis der gehaltenen Qualitätsmaße für die neuronalen Kandidatennetzwerke, neuer Werte der Hyperparameter und der Netzwerkparameter für das neuronale Kandidatennetzwerk Folgendes umfasst:
- Feststellen, ob das aktualisierte Qualitätsmaß besser ist als jedes gehaltene Qualitätsmaß; und  
Setzen, als Reaktion auf die Feststellung, dass das aktualisierte Qualitätsmaß nicht besser ist als jedes gehaltene Qualitätsmaß, der neuen Werte der Hyperparameter auf die gehaltenen Werte der Hyperparameter für das neuronale Kandidatennetzwerk mit dem besten gehaltenen Qualitätsmaß.
11. Verfahren nach einem der Ansprüche 1 bis 10, wobei das Trainieren des neuronalen Kandidatennetzwerks mittels des iterativen Trainingsprozesses und gemäß den gehaltenen Werten der Hyperparameter für das neuronale Kandidatennetzwerk, bis Terminierungskriterien erfüllt sind, das Trainieren des neuronalen Kandidatennetzwerks umfasst, bis eine Schwellenanzahl von Iterationen der iterativen Trainingsprozesse durchgeführt wurde.
12. System, das eine Vielzahl von Recheneinheiten und eine oder mehrere Speichervorrichtungen umfasst, die Befehle speichern, die bei Ausführung durch das System die Vielzahl von Recheneinheiten zum Durchführen der Operationen des jeweiligen Verfahrens nach einem der Ansprüche 1 bis 11 veranlassen.
13. Computerspeichermedium(medien), das (die) Befehle speichert (speichern), die bei Ausführung durch ein eine Vielzahl von Recheneinheiten umfassendes System die Vielzahl von Recheneinheiten zum Durchführen der Operationen des jeweiligen Verfahrens nach einem der Ansprüche 1 bis 11 veranlassen.

### Revendications

1. Procédé de formation d'un réseau neuronal comportant une pluralité de paramètres de réseau pour réaliser une tâche de réseau neuronal particulière et pour déterminer des valeurs formées des paramètres de réseau en utilisant un processus de formation itératif comportant une pluralité d'hyper-paramètres, le procédé étant mis en oeuvre par un système distribué comprenant une pluralité d'unités informatiques, le procédé comprenant :
- le maintien d'une pluralité de réseaux neuronaux candidats (120A, 120N) et, pour chacun des réseaux neuronaux candidats, de données spécifiant : (i) des valeurs respectives des paramètres de réseau (125A, 125N) pour le réseau neuronal candidat, (ii) des valeurs respectives des hyper-paramètres (130A, 130N) pour le réseau neuronal candidat et (iii) une mesure de qualité (135A, 135N) qui mesure une performance du réseau neuronal candidat sur la tâche de réseau neuronal particulière ;  
pour chaque réseau neuronal candidat de la pluralité de réseaux neuronaux candidats (120A, 120N), l'attribution d'une unité informatique respective et la réalisation de façon répétée, par l'unité informatique attribuée, des opérations de formation suivantes, les opérations de formation étant réalisées en parallèle pour chaque réseau neuronal candidat et de façon asynchrone par rapport à la réalisation des opérations de formation pour chaque

autre réseau neuronal candidat :

5 la formation (202) du réseau neuronal candidat en utilisant le processus de formation itératif et conformément aux valeurs maintenues des hyper-paramètres pour le réseau neuronal candidat jusqu'à ce que des critères de fin soient satisfaits pour déterminer des valeurs mises à jour des paramètres de réseau pour le réseau neuronal candidat à partir des valeurs maintenues des paramètres de réseau pour le réseau neuronal candidat,

10 la détermination (204) d'une mesure de qualité mise à jour pour le réseau neuronal candidat conformément aux valeurs mises à jour des paramètres de réseau pour le réseau neuronal candidat et aux valeurs maintenues des hyper-paramètres pour le réseau neuronal candidat,

la détermination (206), sur la base au moins des mesures de qualité maintenues pour les réseaux neuronaux candidats et de la mesure de qualité mise à jour pour le réseau neuronal candidat, de nouvelles valeurs des hyper-paramètres et des paramètres de réseau pour le réseau neuronal candidat,

15 la détermination (208) d'une nouvelle mesure de qualité pour le réseau neuronal candidat conformément aux nouvelles valeurs des paramètres de réseau pour le réseau neuronal candidat et aux nouvelles valeurs des hyper-paramètres pour le réseau neuronal candidat, et

la mise à jour (210) des données maintenues pour le réseau neuronal candidat pour spécifier les nouvelles valeurs des hyper-paramètres, les nouvelles valeurs des paramètres de réseau et la nouvelle mesure de qualité ; et

20 la sélection des valeurs formées des paramètres de réseau parmi les valeurs de paramètre dans les données maintenues sur la base des mesures de qualité maintenues pour les réseaux neuronaux candidats après que les opérations de formation ont été réalisées de façon répétée.

2. Procédé selon la revendication 1, comprenant en outre :

25 la fourniture des valeurs formées des paramètres de réseau pour leur utilisation lors du traitement de nouvelles entrées sur le réseau neuronal.

3. Procédé selon l'une quelconque des revendications 1 ou 2, dans lequel la sélection des valeurs formées des paramètres de réseau parmi les valeurs de paramètre dans les données maintenues sur la base des mesures de qualité maintenues pour les réseaux neuronaux candidats comprend :

30 la sélection des valeurs de paramètre maintenues du réseau neuronal candidat présentant une mesure de qualité maintenue la meilleure de l'un quelconque des réseaux neuronaux candidats après que les opérations de formation ont été réalisées de façon répétée.

4. Procédé selon l'une quelconque des revendications 1 ou 2, dans lequel la sélection des valeurs formées des paramètres de réseau parmi les valeurs de paramètre dans les données maintenues sur la base des mesures de qualité maintenues pour les réseaux neuronaux candidats comprend :

40 la détermination de si la mesure de qualité maintenue pour le réseau neuronal candidat respectif parmi la pluralité de réseaux neuronaux candidats est en-deçà d'un premier pourcentage prédéterminé des réseaux neuronaux candidats ; et

en réponse, l'échantillonnage des valeurs de paramètre maintenues d'un réseau neuronal candidat présentant une mesure de qualité maintenue au-delà d'un deuxième pourcentage prédéterminé des réseaux neuronaux candidats, après que les opérations de formation ont été réalisées de façon répétée.

5. Procédé selon l'une quelconque des revendications 1 à 4, dans lequel la détermination, sur la base au moins des mesures de qualité maintenues pour les réseaux neuronaux candidats, de nouvelles valeurs des hyper-paramètres et des paramètres de réseau pour le réseau neuronal candidat comprend :

50 la détermination de si la mesure de qualité mise à jour est meilleure que chaque mesure de qualité maintenue ; et en réponse à la détermination du fait que la mesure de qualité mise à jour n'est pas meilleure que chaque mesure de qualité maintenue, l'établissement des nouvelles valeurs des paramètres de réseau aux valeurs maintenues des paramètres de réseau pour le réseau neuronal candidat présentant la meilleure mesure de qualité maintenue.

6. Procédé selon la revendication 5, comprenant en outre :

55 en réponse à la détermination du fait que la mesure de qualité mise à jour est meilleure que chaque mesure de qualité maintenue, l'établissement des nouvelles valeurs des paramètres de réseau aux valeurs mises à jour des

paramètres de réseau.

- 5
7. Procédé selon l'une quelconque des revendications 1 à 6, dans lequel la détermination, sur la base au moins des mesures de qualité maintenues pour les réseaux neuronaux candidats, de nouvelles valeurs des hyper-paramètres et des paramètres de réseau pour le réseau neuronal candidat comprend :
- 10
- l'établissement des nouvelles valeurs des hyper-paramètres sur la base des valeurs maintenues des hyper-paramètres pour le réseau neuronal candidat présentant la meilleure mesure de qualité maintenue.
8. Procédé selon la revendication 7, dans lequel l'établissement des nouvelles valeurs des hyper-paramètres comprend la permutation aléatoire d'une ou de plusieurs des valeurs maintenues des hyper-paramètres du réseau neuronal candidat présentant la meilleure mesure de qualité maintenue.
- 15
9. Procédé selon l'une quelconque des revendications 1 à 6, dans lequel la détermination, sur la base au moins des mesures de qualité maintenues pour les réseaux neuronaux candidats, de nouvelles valeurs des hyper-paramètres et des paramètres de réseau pour le réseau neuronal candidat comprend :
- 20
- l'échantillonnage des nouvelles valeurs des hyper-paramètres à partir d'une distribution antérieure sur des valeurs possibles pour les hyper-paramètres ; et/ou  
l'échantillonnage des nouvelles valeurs des paramètres de réseau à partir d'une distribution antérieure sur des valeurs possibles pour les paramètres de réseau.
- 25
10. Procédé selon l'une quelconque des revendications 1 à 4, dans lequel la détermination, sur la base au moins des mesures de qualité maintenues pour les réseaux neuronaux candidats, de nouvelles valeurs des hyper-paramètres et des paramètres de réseau pour le réseau neuronal candidat comprend :
- 30
- la détermination de si la mesure de qualité mise à jour est meilleure que chaque mesure de qualité maintenue ; et en réponse à la détermination du fait que la mesure de qualité mise à jour n'est pas meilleure que chaque mesure de qualité maintenue, l'établissement des nouvelles valeurs des hyper-paramètres aux valeurs maintenues des hyper-paramètres pour le réseau neuronal candidat présentant la meilleure mesure de qualité maintenue.
- 35
11. Procédé selon l'une quelconque des revendications 1 à 10, dans lequel la formation du réseau neuronal candidat en utilisant le processus de formation itératif et conformément aux valeurs maintenues des hyper-paramètres pour le réseau neuronal candidat jusqu'à ce que des critères de fin soient satisfaits comprend la formation du réseau neuronal candidat jusqu'à ce qu'un nombre seuil d'itérations du processus de formation itératif aient été réalisées.
- 40
12. Système comprenant une pluralité d'unités informatiques et un ou plusieurs dispositifs de stockage stockant des instructions qui, lorsqu'elles sont exécutées par le système, amènent la pluralité d'unités informatiques à réaliser les opérations du procédé respectif selon l'une quelconque des revendications 1 à 11.
- 45
13. Support(s) de stockage informatique stockant des instructions qui, lorsqu'elles sont exécutées par un système comprenant une pluralité d'unités informatiques, amènent la pluralité d'unités informatiques à réaliser les opérations du procédé respectif selon l'une quelconque des revendications 1 à 11.
- 50
- 55

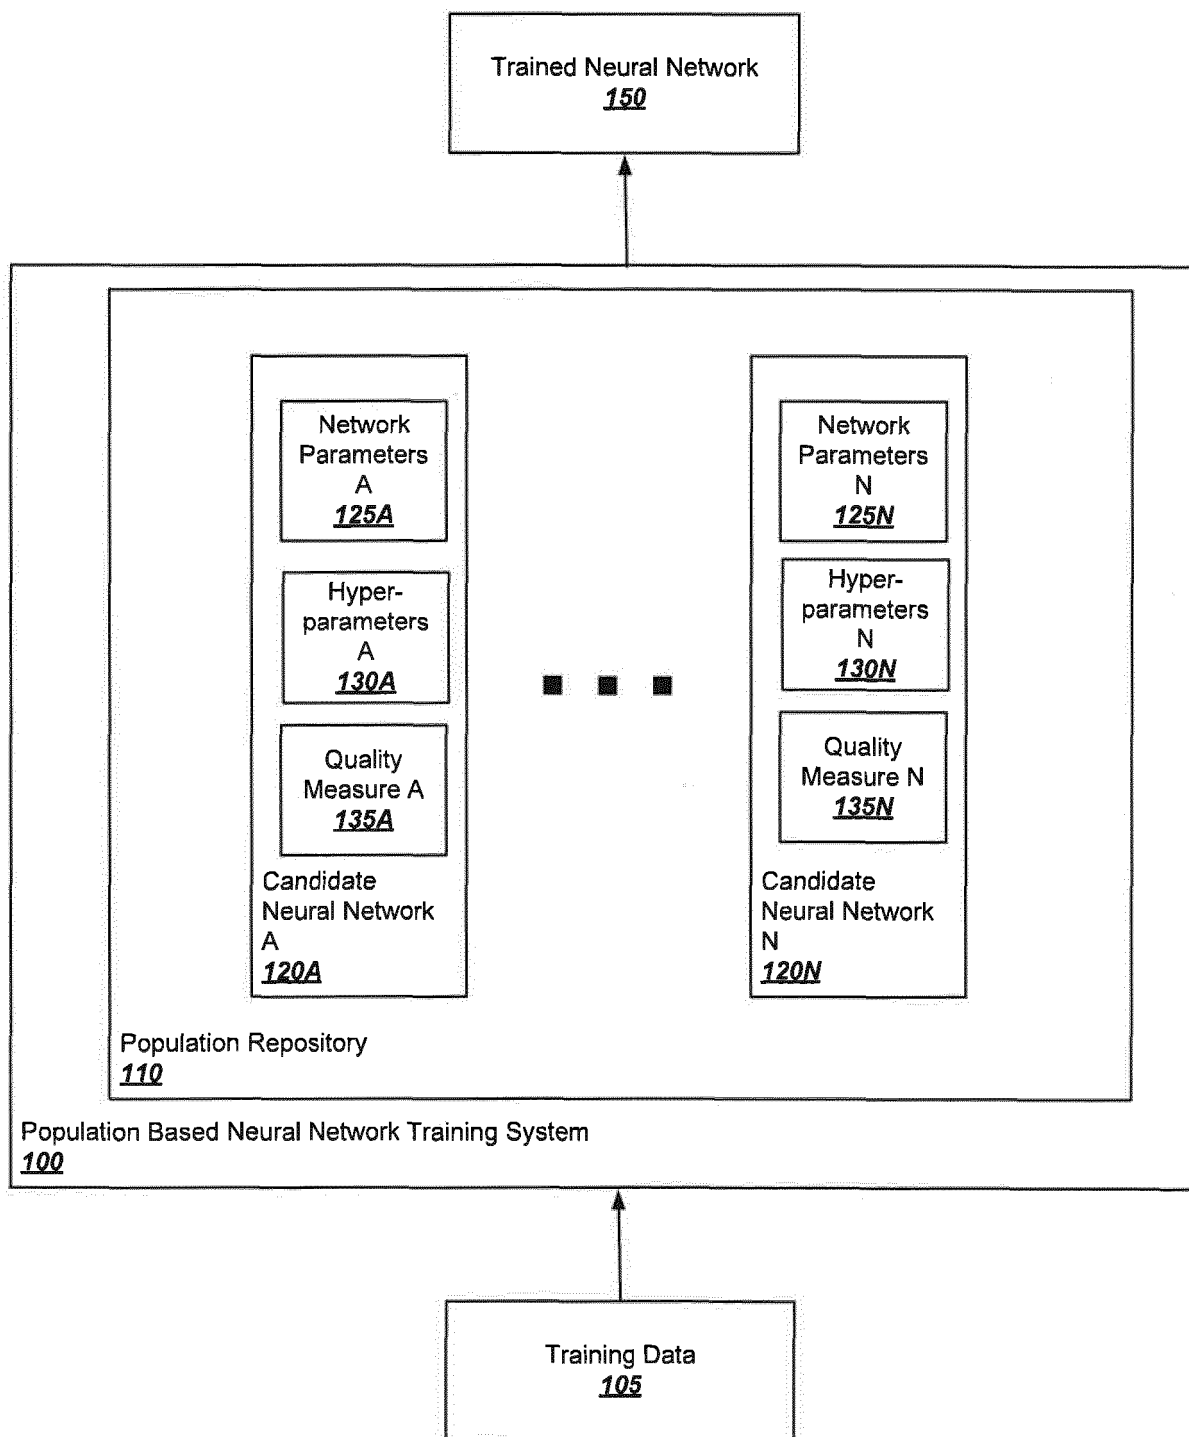


FIG. 1

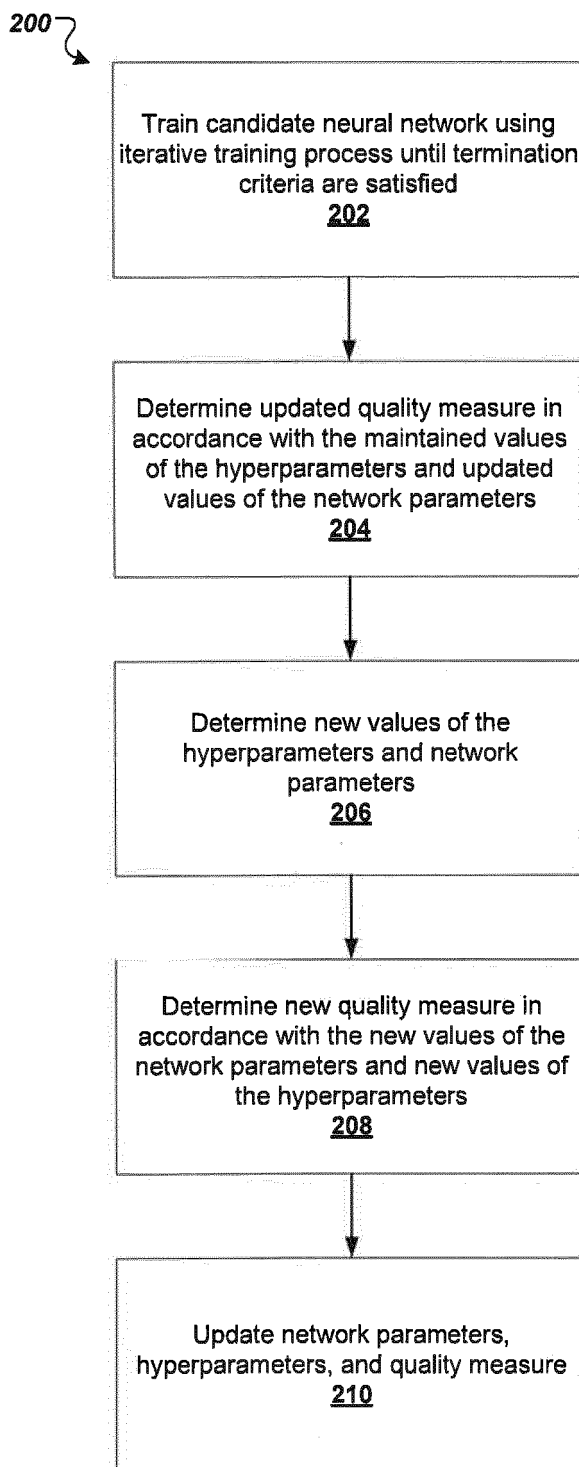
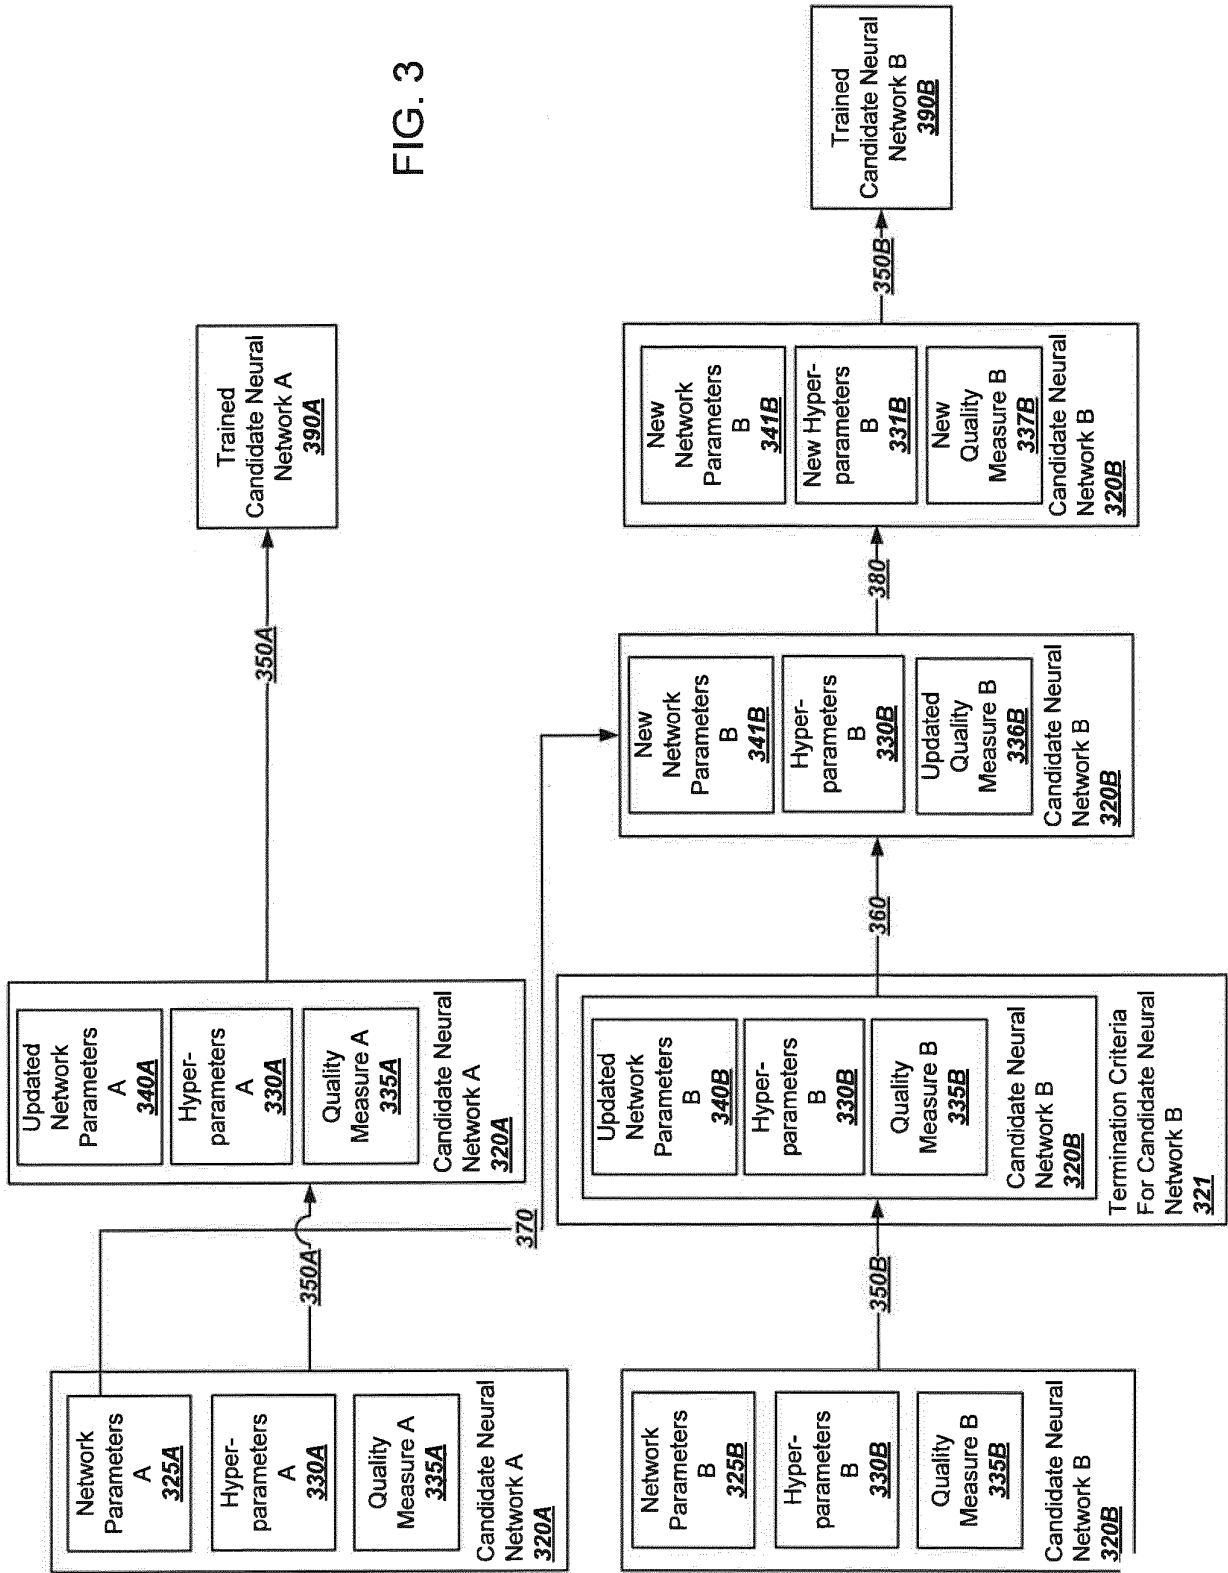


FIG. 2

FIG. 3



**REFERENCES CITED IN THE DESCRIPTION**

*This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.*

**Patent documents cited in the description**

- WO 62590177 A [0001]

**Non-patent literature cited in the description**

- **FERNANDO**. Convolution by Evolution: Differentiable Pattern Producing Networks. *arXiv:1606.02580v1*, 08 June 2016 [0003]